Ex/CSE/PE/H/T/423A/2025

B.E. COMPUTER SCIENCE & ENGINEERING 4th YEAR 2nd SEMESTER EXAM- 2025

DISTRIBUTED COMPUTING

Time: 3 hrs Full Marks: 100

Group A (Total Marks: 30) [CO1]
Answer Question No. 1 (Compulsory) and Question No. 2 OR Question No. 3

1. a) What are the different motivating factors which are important for building an application over a distributed system? Consider two distributed applications i) data sharing and collaboration ii) E- commerce and explain your choices with proper justification.
   b) What are the different techniques to handle a single point of failure? Explain how SPOF can be managed in following applications/services i) Database server ii) Network switch iii) ISP
   (5+5)+(5+5)=20
2. Explain how a distributed system projects a single system image. What is the role of transparency in a distributed computing environment? What are the different forms of transparency? 4+2+4=10
3. What are the benefits of distributed systems? Discuss some disadvantages of distributed systems that local only systems do not show. Compare redundancy, load balancing, and failover mechanisms in respect to performance and cost. 5+5=10

Group B (Total Marks: 30) [CO2 and CO3]
Answer Question No. 4 (Compulsory) and Question No. 5 OR Question No. 6

4. a) Explain the working procedure of BitTorrent with a suitable diagram. Describe the AWS serverless platforms and its benefits with Compute, Storage, Data Stores and API Proxy services.
   b) Describe the peer selection approaches used in Napster, Gnutella and KaZaA. Explain the working procedure of Spark with a suitable example and diagram.
   (5+5)+(5+5)=20
5. Explain different functions which are essential to achieve high availability in a cluster computing environment. Describe the basic architecture and main components of Resource Management and Scheduling (RMS). Compare and contrast cloud computing with grid computing. 3+4+3=10
6. Describe pre-copy and post-copy VM live migration techniques. Compare these two techniques in terms of performance metrics. Explain the split driver model based I/O virtualization technique in XEN with a suitable diagram. 6+4=10

[Turn over

Group C (Total Marks: 40) [CO4]
Answer Question No. 7 (Compulsory) and Question No. 8 OR Question No. 9

7. a) In the symmetric mode of synchronization in NTP, suppose you are given that server A and server B are connected by a symmetric channel, i.e., the channel shows the same (but unknown) message delay both ways, i.e., the A to B delays is the same as B to A delay. Show using the equations for analyzing the NTP protocol, that under this situation, one can estimate the clock skew accurately (i.e., with an error of 0).

b) Give an example where the Lamport's clock algorithm comes short (i.e., the Lamport's algorithm cannot clearly conclude that event e happens before e', even those L(e) < L(e'), where L(e) is the Lamport's timestamp of the event e), and the vector clock algorithm concludes clearly that event e happened before e' or NOT.

c) Consider the following figure showing three processes P1, P2, P3.

i) Is the run a linearization of events? Explain why or why not.

ii) Is the cut, shown by curve X, a consistent cut? Why?

iii) Determine two other consistent cuts in Figure 2 and specify their frontiers.

iv) Determine all the events that happen before event e_2^4 (as per Lamport's Happened-Before relation). Also determine the events that are concurrent with event e_2^4.

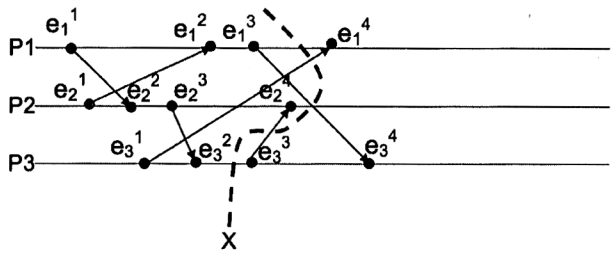

5+5+10=20

8. a) You are synchronizing your clock from a time server using Cristian's algorithm and observe the following times:

- timestamp at client when the message leaves the client: 6:22:15.100
- timestamp generated by the server: 6:21:10.700
- timestamp at client when the message is received at client: 6:22:15.250

To what value do you set the client's clock? If the best-case *round-trip* message transit time is 124 msec (0.124 sec), what is the error of the clock on the client?

b) A system of four processes, (P_1, P_2, P_3, P_4), performs the following events:

a. P1 sends a message to P3 (to event e).
b. P_1 receives a message from P_3 (from event g).
c. P_2 executes a local event.
d. P_2 receives a message from P_3 (from event f).
e. P_3 receives a message from P_1 (from event a).
f. P_3 sends a message to P_2 (to event d).
g. P_3 sends a message to P_1 (to event b).
h. P_4 executes a local event.

When taking place on the same processor, the events occur in the order listed. Assign Lamport timestamps to each event. Assume that the clock on each processor is initialized to 0 and incremented before each event. For example, event *a* will be assigned a timestamp of 1.

a. **1** b. c. d.

e. f. g. h.

c) Assign vector timestamps to each event (above). Assume that the vector clock on each processor is initialized to (0,0,0,0) with the elements corresponding to (P_1, P_2, P_3, P_4). For example, event *a* will be assigned a timestamp of (1, 0, 0, 0). Which events are concurrent with event *d*?

d) Event a has a Lamport timestamp of 4. Event b has a Lamport timestamp of 8. What can we tell about events a and b? You are on a network with a 9 Mbps downstream bandwidth and a 1 Mbps upstream bandwidth. Your client makes a request at 9:10:00.0 and gets a response 200 msec later. If the time on the server is 8:42:00.0, what would the time be set to on your client?

5+5+5+5=20

9. a) Is the following data store sequentially consistent? Justify your answer. Explain weak consistency with a valid and invalid sequence of events. Explain the working procedure of MapReduce to scale data processing over multiple computing nodes with a suitable example and diagram.

A	W(x)a		W(x)b	
B		R(x)a	R(x)b	
C			R(x)b	R(x)a

b) Explain the concept of the "Release consistency model". What modification over "Release consistency model" leads to "Lazy release consistency model"? Explain with a proper example and diagram.

(3+3+4)+(5+5)=20